Optimal control of path-dependent McKean-Vlasov SDEs in infinite dimension

Idris Kharroubi
joint work with
A. Cosso, F. Gozzi, H. Pham & M. Rosestolato

9th International Colloquium on BSDEs and Mean Field Systems
June 27-July 01, 2022
Annecy - France

Outline

Motivation

- Optimal control of McKean-Vlasov type SDEs is a recent topic which arises in a natural way in problems where the dynamics of the state equation depends on the state/control law.
- Typical "planner's problem" in economics with many players, as opposed to "agent's problem" ($\longrightarrow$ Mean Field Game).
- Various papers recently have studied the case when state equation is finite dimensional (see e.g. the book [Carmona-Delarue] and various papers) and possibly path-dependent [Wu-Zhang].

- Up to our knowledge, no paper studies the infinite dimensional case. However such case arises naturally in applications.
- We then aim to develop the theory in this case trying also to clarify some issues left in previous papers. This paper is the first step.

Example 1: spatial economic growth models

Spatial growth models (see e.g. [Boucekkine-Camacho-Fabbri '16] [Gozzi-Leocata '21]) lead to optimal control in an infinite dimensional space H.

Typical issue in such problems: the productivity depends on the average of the capital distribution (see e.g. [Turnovsky '06]). $\Rightarrow$ state equation for the capital trajectory $k(\cdot)$ as

$$dk(t) \;=\; [Ak(t) + a(\mathbb{E}k(t))k(t) - \delta k(t) - c(t)]\,dt + \sigma(k(t))dB(t).$$

Here A is the laplace operator, $c(\cdot)$ is the control (consumption rate), $a(\cdot)$, δ, $\sigma(\cdot)$ are given data.

Also need to include in such type of models, delay/path-dependent features like time-to build or vintage capital.

Example 2: Lifecycle portfolio with "sticky" wages

In such problems (see e.g. [Djeiche-Gozzi-Zanco-Zanella '20]), labor income "$y(\cdot)$" (one of the state equations of the optimal portfolio problem) described by one-dimensional delay SDEs of McKean-Vlasov type as follows

$$dy(t) = \left[b_0(\mathbb{P}_{y(t)}) + \int_{-d}^{0} y(t+\xi)\,\phi(\xi)\,d\xi \right] dt + \sigma\, y(t)\, dZ(t).$$

(here $\phi \in L^2(-d,0;\mathbb{R})$ is a given datum and Z is a one-dimensional Brownian motion). Such equations can be rephrased as SDEs in the Hilbert space $\mathbb{R} \times L^2(-d,0;\mathbb{R})$ and the resulting dynamics is a Mc Kean - Vlasov SDEs in infinite dimension.

Outline

Basic and probabilistic setting

- The state space H and the control space K are real separable Hilbert spaces. The control set U is a Polish space or a Borel subset of it. The horizon T is finite.
- $(\Omega, \mathscr{F}, \mathbb{P}, (\mathscr{F}_s^B)_{s\in[0,T]}, B)$ is a reference probability space i.e. a complete probability space with a cylindrical Brownian motion with values in K and $(\mathscr{F}_s^B)_{s\in[0,T]}$ is the augmented filtration generated by B.
- There exists a sub-σ-algebra $\mathscr{G}$ of $\mathscr{F}$, independent of $\mathscr{F}_\infty^B$ and admitting a $\mathscr{G}$-measurable r.v. $U_{\mathscr{G}}$ with uniform distribution in $[0,1]$.
- We call $\mathbb{F} := (\mathscr{F}_s)_{s\in[0,T]}$ with $\mathscr{F}_s = \mathscr{G} \vee \mathscr{F}_s^B$,

State equation

$X(\cdot)$ is the state while $\alpha(\cdot)$ is the control:

$$\begin{aligned} dX(s) &= AX(s) + b(s, X_s, \mathbb{P}_{X_s}, \alpha(s), \mathbb{P}_{\alpha(s)})ds \\ &\quad + \sigma(s, X_s, \mathbb{P}_{X_s}, \alpha(s), \mathbb{P}_{\alpha(s)})dB(s), \quad s > t \end{aligned}$$

and $X(s) = \xi(s)$ for $0 \leq s \leq t$. Here:

- A is a linear unbounded operator (generating a strongly continuous semigroup);
- b and σ are progressive;
- subscript $_s$ denotes the history of the process up to time s.

Assume

- b and σ are bounded and satisfy a Lipschitz condition wrt state and measure (in the Wasserstein space $(\mathscr{P}_2, \mathscr{W}_2)$);
- the initial datum ξ belongs to $\mathbf{S_2}(\mathbb{F})$ which is the set of H-valued continuous $\mathbb{F}$-progressively measurable processes ξ such that $\|\xi\|_{S_2} := \left(\mathbb{E}[\sup_{t\in[0,T]} |\xi(t)|_H^2]\right)^{1/2} < +\infty$.

Then, the controlled process X is well defined in $\mathbf{S_2}(\mathbb{F})$.

Objective function

We aim to maximize the functional

$$J(t,\xi;\alpha) = \mathbb{E}\left[g\left(X_T^{t,\xi,\alpha},\mathbb{P}_{X_T^{t,\xi,\alpha}}\right) + \int_t^T f\left(s, X_s^{t,\xi,\alpha}, \mathbb{P}_{X_s^{t,\xi,\alpha}}, \alpha(s), \mathbb{P}_{\alpha(s)}\right) ds\right]$$

Here:

- f and g are progressive, continuous, locally bounded and locally uniformly continuous in state and measure (uniformly wrt the other variables) and with quadratic growth in state.
- $\alpha \in \mathscr{U}$ where $\mathscr{U}$ is the space of $\mathbb{F}$-progressively measurable processes $[0,T] \times \Omega \to U$.
- $X^{t,\xi,\alpha}$ is the solution of the state equation with initial data $(t,\xi) \in [0,T] \times \mathbf{S}_2(\mathbb{F})$ and control $\alpha \in \mathscr{U}$.

Value function

The value function $V : [0, T] \times \mathbf{S}_2 \to \mathbb{R}$ is defined by

$$V(t,\xi) = \sup_{\alpha(\cdot)\in\mathscr{U}} J(t,\xi;\alpha), \qquad (t,\xi) \in [0, T] \times \mathbf{S}_2 \,.$$

On the line of the previous papers on the finite dimensional case, in particular [Wu-Zhang,'20], we expect that the value function only depends on the law of the initial datum

We also expect that its equivalent v on measures solves, in a suitable viscosity sense, the associated HJB equation on the space of measures $\mathscr{P}_2(C([0, T]; H))$.

Outline

Summary of the results

In the paper we prove the following.

- Law invariance of V, i.e. that $V(t,\xi)$ only depends on the law of ξ up to time t.
- Dynamic Programming Principle.
- Ito formula in the context of pathwise derivatives in the spirit of Dupire and Wu-Zhang definition.
- Viscosity property of V for the HJB equation.
- Uniqueness in a specific case.

Comments of results

Obtained results are those expected (as the community becomes more and more used to this mean field framework).

However, proofs involve a lot of technicalities (the paper is 54 pages long!).

Rest of the talk : try to stay reasonable (focus on the main steps)

Some comments on further ongoing work will also be given.

Law invariance of V

- The goal is to prove that, if ξ and η belong to $\mathbf{S}_2(\mathbb{F})$, with $\mathbb{P}_\xi = \mathbb{P}_\eta$, then $V(t,\xi) = V(t,\eta)$
- This was proved, in the finite dimensional case, in [Cosso-Pham,'18]. However their proof uses a result from [Aliprantis-Border,'06], Corollary 18.23, which is not correct as it is. Hence their proof does not work.
- Our proof is based on the fact that one can find, for every ξ, η above, two r.v. U_ξ and U_η, with uniform law on $[0,1]$, such that ξ and U_ξ (and also η and U_η) are independent. ⚠
- We also provide and example where this does not apply.

Define v by $v(t,\mu) = V(t,\xi)$ for any ξ with $\mathbb{P}_\xi = \mu$ (lift-inverse).

DPP for V

Theorem

Under Assumption on the coefficients b, σ, f, g, the lifted value function V satisfies the ***dynamic programming principle****:*

$$V(t,\xi) = \sup_{\alpha \in \mathcal{U}} \left\{ \mathbb{E}\left[\int_t^s f_r(X^{t,\xi,\alpha}, \mathbb{P}_{X^{t,\xi,\alpha}_{\cdot\wedge r}}, \alpha_r, \mathbb{P}_{\alpha_r}) \, dr \right] + V(s, X^{t,\xi,\alpha}) \right\}$$

for every $t, s \in [0,T]$, with $t \leq s$, and $\xi \in \mathbf{S}_2(\mathbb{F})$.

Rq: no measurable selection issue as the function V depends on the whole r.v. ξ.

DPP for the lift inverse v

Corollary

*Under previous assumptions on the coefficients b, σ, f, g, the value function v satisfies the **dynamic programming principle**:*

$$v(t,\mu) = \sup_{\alpha \in \mathcal{U}} \left\{ \mathbb{E}\left[\int_t^s f_r(X^{t,\xi,\alpha}, \mathbb{P}_{X^{t,\xi,\alpha}_{\cdot\wedge r}}, \alpha_r, \mathbb{P}_{\alpha_r})\, dr \right] + v(s, \mathbb{P}_{X^{t,\xi,\alpha}}) \right\},$$

for every $t, s \in [0,T]$, with $t \leq s$, $\mu \in \mathscr{P}_2(C([0,T];H))$ and $\xi \in \mathbf{S}_2(\mathbb{F})$ with $\mathbb{P}_\xi = \mu$.

Since V non-anticipative, namely $V(t,\xi) = V(t, \xi_{\cdot\wedge t})$, for every $(t,\xi) \in [0,T] \times \mathbf{S}_2(\mathbb{F})$, v satisfies

$$v(t,\mu) = v(t, \mu_{[0,t]}),$$

where $\mu_{[0,t]} = \mu \circ \left((x_s)_{s\in[0,T]} \mapsto (x_{s\wedge t})_{s\in[0,T]} \right)^{-1}$.

Derivatives

We need to define derivatives in this path-dependent framework. That is

- time derivative,
- measure derivative,
- path$\times$measure derivative.

Use a definition via the lift function. Denote by Φ the lift of φ:

$$\Phi(t,\xi) = \varphi(t,\mathbb{P}_\xi), \qquad \forall\,(t,\xi).$$

Time derivative

Define

$$\mathscr{H} = [0,T] \times \mathscr{P}_2\big(C([0,T];H)\big)\,.$$

φ is pathwise differentiable in time at $(t,\mu) \in \mathscr{H}$, if the following limit

$$\partial_t \varphi(t,\mu) \quad = \quad \lim_{\delta \to 0^+} \frac{\varphi(t+\delta, \mu_{[0,t]}) - \varphi(t,\mu)}{\delta}$$

exists and is finite.

Measure derivative

Φ is **pathwise differentiable in space** at (t,ξ) if there exists $D\Phi(t,\xi) \in L^2(\Omega;H)$ such that

$$\lim_{Y\to 0} \frac{\left|\Phi(t,\xi + Y\mathbf{1}_{[t,T]}) - \Phi(t,\xi) - \mathbb{E}\big[\langle D\Phi(t,\xi), Y\rangle_H\big]\right|}{|Y|_{L^2(\Omega;H)}} = 0.$$

Then, there exists a measurable function $\partial_\mu\varphi(t,\mu)(.)$ such that

$$D\Phi(t,\xi) = \partial_\mu\varphi(t,\mu)(\xi) \qquad \mathbb{P}\text{-a.s.}$$

The function $\partial_\mu\varphi(t,\mu)$ is the measure derivative of φ at $(t,\mu) \in \mathscr{H}$.

Path × measure derivative

We say that φ is **pathwise differentiable in measure and space at** (t,μ,x) if there exists an operator $\partial_x\partial_\mu\varphi(t,\mu)(x)\in\mathscr{L}(H)$ such that

$$\lim_{h\to 0}\frac{\left|\partial_\mu\varphi(t,\mu)(x+h\mathbf{1}_{[t,T]})-\partial_\mu\varphi(t,\mu)(x)-\partial_x\partial_\mu\varphi(t,\mu)(x)h\right|_H}{|h|_H}=0.$$

We refer to $\partial_x\partial_\mu\varphi(t,\mu)(x)$ as the **second-order pathwise derivative in measure and space of** φ **at** (t,μ,x).

$\boldsymbol{C}_b^{1,2}(\mathscr{H})$: φ s.t. φ, $\partial_t\varphi$, $\partial_\mu\varphi$, $\partial_x\partial_\mu\varphi$ are cont. and bounded.

Fix $t \in [0,T]$ and $\xi \in \mathbf{S}_2(\mathbb{F})$.

Let $F\colon [0,T]\times\Omega \to H$, $G\colon [0,T]\times\Omega \to \mathscr{L}_2(K;H)$ $\mathbb{F}$-progressive, such that

$$\int_0^T \mathbb{E}[|F_s|_H^2]\,ds < \infty, \qquad \int_0^T \mathbb{E}\big[\mathrm{Tr}\big(G_s G_s^*\big)\big]\,ds < \infty.$$

Consider the process $X = (X_s)_{s\in[0,T]}$ given by

$$X_s = \xi_{s\wedge t} + \int_t^{s\vee t} F_r\,dr + \int_t^{s\vee t} G_r\,dB_r, \qquad \forall\, s \in [0,T].$$

Theorem

If $\varphi : \mathscr{H} \to \mathbb{R}$ is in $\mathbf{C}_b^{1,2}(\mathscr{H})$, then the following ***Itô formula*** *holds:*

$$
\begin{aligned}
\varphi(s, \mathbb{P}_{X_{\cdot\wedge s}}) \quad = \quad & \varphi(t, \mathbb{P}_{\xi_{\cdot\wedge t}}) + \int_t^s \partial_t \varphi(r, \mathbb{P}_{X_{\cdot\wedge r}})\, dr \\
& + \int_t^s \mathbb{E}\big[\langle F_r, \partial_\mu \varphi(r, \mathbb{P}_{X_{\cdot\wedge r}})(X_{\cdot\wedge r})\rangle_H\big]\, dr \\
& + \frac{1}{2} \int_t^s \mathbb{E}\big[\mathrm{Tr}\big(G_r G_r^* \partial_x \partial_\mu \varphi(r, \mathbb{P}_{X_{\cdot\wedge r}})(X_{\cdot\wedge r})\big)\big]\, dr,
\end{aligned}
$$

for every $s \in [t, T]$.

HJB equation

Let $\mathscr{M}_t := \{\mathfrak{a} \colon \Omega \to U \colon \mathfrak{a} \text{ is } \mathscr{F}_t\text{-measurable}\}$. HJB PDE writes

$$
\begin{aligned}
0 \;=\; & \partial_t w(t,\mu) + \mathbb{E}\langle \xi_t, A^* \partial_\mu w(t,\mu)(\xi)\rangle_H \\
& + \sup_{\mathfrak{a}\in\mathscr{M}_t} \Big\{ \mathbb{E}\big[f_t(\xi,\mu,\mathfrak{a},\mathbb{P}_{\mathfrak{a}}) + \big\langle b_t(\xi,\mu,\mathfrak{a},\mathbb{P}_{\mathfrak{a}}), \partial_\mu w(t,\mu)(\xi)\big\rangle_H \big] \\
& + \frac{1}{2}\mathbb{E}\Big[\mathrm{Tr}\Big(\sigma_t(\xi,\mu,\mathfrak{a},\mathbb{P}_{\mathfrak{a}})\sigma_t^*(\xi,\mu,\mathfrak{a},\mathbb{P}_{\mathfrak{a}}) \partial_x \partial_\mu w(t,\mu)(\xi)\Big)\Big] \Big\},
\end{aligned}
$$

for $(t,\mu) \in \mathscr{H}$, $t < T$, $\xi \in \mathbf{S}_2(\mathscr{G})$ s.t. $\mathbb{P}_\xi = \mu$,

with terminal condition

$$
w(T,\mu) \;=\; \mathbb{E}[g(\xi,\mu)]
$$

for $\mu \in \mathscr{P}_2(C([0,T];H))$, $\xi \in \mathbf{S}_2(\mathscr{G})$ s.t. $\mathbb{P}_\xi = \mu$.

Regular solutions

- We say that a function $w\colon \mathscr{H} \to \mathbb{R}$ belongs to the space $\boldsymbol{C}^{1,2}_{b,A^*}(\mathscr{H})$ if it satisfies the following regularity assumptions:

(i) $w\colon \mathscr{H} \to \mathbb{R}$ belongs to $\boldsymbol{C}^{1,2}_{b}(\mathscr{H})$;

(ii) for all $(t,\mu,\xi) \in H \times \mathbf{S}_2(\mathbb{F})$, $\partial_\mu \varphi(t,\mu)(\xi) \in L^2(\Omega; D(A^*))$ and the map

$$\mathscr{H} \times \mathbf{S}_2(\mathbb{F}) \longrightarrow L^2(\Omega; H), \qquad (t,\mu,\xi) \longmapsto A^*\varphi(t,\mu)(\xi)$$

is continuous and bounded.

- We say that a function $w\colon \mathscr{H} \to \mathbb{R}$ is a classical solution to the HJB equation if it belongs to the space $\boldsymbol{C}^{1,2}_{b,A^*}(\mathscr{H})$ and satisfies the HJB PDE.

Regular solutions

Theorem

Suppose that previous assumptions on b, σ, f, g hold and that b, σ, f are uniformly continuous in t, uniformly with respect to the other variables. Assume that the value function v belongs to the space $\boldsymbol{C}^{1,2}_{b,A^*}(\mathscr{H})$. Then v is a classical solution of the HJB PDE.

Consequence of Itô formula and DPP applied to the function v.

viscosity solutions

Definition of viscosity solutions same as classical one except that it uses test functions $\varphi \in \boldsymbol{C}^{1,2}_{b,A^*}(\mathscr{H})$.

Theorem

Under previous assumptions on b, σ, f and g, the value function v is a viscosity solution to the HJB PDE.

The proof follows exactly the same lines as for the regular case, simply replacing v with a test function φ.

Comparison: back to $\mathbb{R}^d$ and no path dependency

Second-order HJB equations in the Wasserstein space is emerging and still at an early stage.

[Burzoni et al. 20] study a special class of HJB equations: b, σ, f, g do not depend on x, and control is deterministic.

[Wu & Zhang 20] use a notion of viscosity solution withconditions formulated on compact subsets of the Wasserstein space.

Comparison

Theorem (Comparison)

Suppose that previous assumptions hold, σ does not dependent on μ (the law) and is C_b^2.
Let $u_1, u_2 : [0,T]\times\mathscr{P}_2(\mathbb{R}^d)\to\mathbb{R}$ be continuous bounded functions. Suppose that u_1 (resp. u_2) is a viscosity subsolution (resp. supersolution) of HJB equation . Then

$$u_1 \quad \leq \quad u_2, \qquad \text{on } [0,T]\times\mathscr{P}_2(\mathbb{R}^d).$$

Corollary (Uniqueness)

Under previous assumptions, v is the unique bounded and continuous viscosity solution of HJB equation.

Not possible to use classical Ishii's Lemma based approach.

Back to the original proof: prove that $u_1 \leq v \leq u_2$.

Issues

- Need a regular version/approximation of $v \Rightarrow$ replace v by its n-agent/particle (common) optimization value v_n.
- No local compactness $\Rightarrow$ use a smooth variational principle of Borwein-Preiss type with perturbation/gauge function constructed via a partition of the space as in [Dereich et al.13] and a convolution with Gaussians.

Outline

- Cases where the solutions of HJB equation are regular and the optimal feedback control can be found.
- Mean Field Games in infinite dimension (with S. Federico and M. Rosestolato).
- Applications to specific problems.

Thank you for your attention